

Remarks

The amendment to claim 11

5 The amendment removes an “is” that was erroneously not deleted when claim 11 was amended in response to the final Office action of 3/24/04. For the reasons set forth in Applicants’ response to the final Office action mailed 4/5/04, the claim as amended is fully supported by the Specification as filed.

10 Traversal of the rejections

Explanation of the invention: FIGs. 4 and 5 and page 19, line 16-page 21, line 5

A problem in any distributed system is making sure that changes in the components of the digital systems do not result in inconsistencies in the distributed system. The transactions that make the changes often employ protocols such as the two-phase commit protocol to ensure 15 consistency. These protocols represent system overhead, i.e., they require the use of resources in the distributed system that would otherwise be available to users of the distributed system. For that reason, various techniques have been developed for optimizing the protocols.

The invention is a novel technique for optimizing such protocols. The technique is based on the notion that the overhead required for a protocol can be reduced if the component of the system 20 that is managing the protocol has available to it constantly-updated information that is relevant to the protocol and uses that information to optimize the protocol.

In the exemplary embodiment, the distributed system is a distributed database system and the components are component database systems. The two-phase commit protocol is used to ensure that when a transaction involving copies of writable data in the component database systems is 25 finished, all of the copies of the writable data are consistent. A characteristic of many distributed database systems is that they include some components in which data involved in a given transaction is read only and others in which it is writable. If the component database system that is managing the transaction knows which components are writable with regard to the transaction, the managing database system can optimize the two-phase commit protocol with 30 regard to the read only components. Since the read only property of a component with regard to

a component changes dynamically, what is needed is a technique which permits information as to the status of the component to be constantly returned to the managing database system.

As shown in FIG. 4, the technique has two components: augmented messages and retained state. The messages that are sent from component systems to the managing system during the

5 transaction are *augmented* with information that indicates whether the component system sending the message is currently read-only with regard to the transaction. The *augmented message*, shown in FIG. 4 at 401, consists of a message that would be normally sent from the component system to the managing system during the transaction (403) to which cohort status information 405 has been added. Included in cohort status information 405 is a read only flag
10 407 which indicates whether the sending component system is currently read only with regard to the transaction.

The *retained state* is state for the transaction in the managing system which is updated from each augmented message for the transaction. In a preferred embodiment, the managing system maintains an outgoing link-transaction object 411 for each transaction for each of the

15 component systems that are involved in a transaction. Included in a link 411(i) for a given transaction and component system is current cohort status 413(i), which embodies the retained state. Included in this state is a read only status flag 415(i), which is set from read-only status flag 407 each time a message for the link's transaction is received from the link's component system. The managing system uses the retained state to optimize the two-phase commit
20 protocol.

Fig. 5 shows how this is done. At the beginning of the transaction, the managing system sets status flag 415(i) in the outgoing link transaction object to *read only* (505). Then, until the transaction is terminated (507), loop 513 is executed, in which the managing system receives each augmented message from the link transaction object's component system (515) and sets

25 cohort status information 413 from the cohort status information 405 in the augmented message. On termination of the transaction (509), the managing system reads the current cohort status 413 for the transaction (519). If it is read only, the managing system optimizes the two-phase commit protocol by only sending the protocol's *abort* message to the component system (527); otherwise, it does the entire two-phase commit protocol with the cohort (529).

30 *The independent claims*

The independent claims claim the optimization technique both from the point of view of the managing system and from the point of view of the component system. Claim 5 is an example of a claim written from the point of view of the component system:

5 **5. (previously presented)** A method of ensuring that a first component of a distributed system that exchanges messages that belong to a transaction with one or more other components of the distributed system is additionally aware of a state of an other component, the state being relevant to a protocol that is employed with the transaction and
the method comprising the steps practiced in the other component of:
10 determining the relevant state; and
 augmenting a message of the transaction with state information indicating the relevant state of the other component,
 the first component using the state information to optimize the protocol.

15 The claimed method's purpose is to "ensur[e] that a first component of a distributed system that exchanges messages that belong to a transaction with one or more other components of the distributed system is additionally aware of a state of an other component . . . [that is] relevant to a protocol . . .". The purpose is achieved by two steps practiced "in the other component": "determining the relevant state" and "augmenting a message of the transaction . . .". The first component is then able to "us[e] the state information to optimize the protocol."

20 Claim 11 is an example of a claim written from the point of view of the managing system:

25 **11. (currently amended)** A method of ensuring that a first component of a distributed system that exchanges messages belonging to a transaction with one or more other components of the distributed system is additionally aware of a state of an other component, the state being relevant to a protocol employed with the transaction and
the method comprising the steps practiced in the first component of:
30 receiving an augmented one of the messages, the augmented message having been augmented by the other component to additionally contain information indicating the relevant state of the other component;
 retaining the relevant state from the augmented message; and
 using the retained relevant state to optimize the protocol.

35 The preamble is like that of claim 5, except that the method steps are "practiced in the first component" and the method steps are "receiving an augmented one of the messages" , "retaining the relevant state from the augmented message", and "using the retained relevant state to optimize the protocol".

The "augmented message" and "retained relevant state" terminology used in the claims has its origins in the Specification as filed. See for example, page 22, lines 1-7:

5

With protocols for transactions, the optimization technique can be used in any situation where the protocol can be optimized if global or local coordinators for the transaction have current knowledge during the transaction of state in the cohorts that is relevant to the optimization. The only requirements for the optimization are that the messages sent from the cohorts in the course of the transaction are augmented with information concerning the relevant state and that the relevant state is available to the coordinator while the coordinator is performing the transaction.

As already indicated above, the "augmented message" appears at 401 in FIG. 4 and the 10 "retained relevant state" appears at 413.

The rejection of the independent claims as anticipated by Burns

What Burns discloses is well described in the patent's *Abstract*:

15 A computing system provides continuous availability of data files that are maintained at a file management system and linked to a database management system (DBMS) through a Datalink data type, even while any particular file is being changed with DBMS append or update operations. When a file is linked, it is designated to be available for read-only operations. A user who wants to perform updates on a file gets a "check-out" copy of the file for updating operations, such that the original file remains linked to the database system while 20 the copy is being updated and remains available to other users. The file management system includes a "check-in" function that receives the updated file, saves the updated file under a new name different from the original, updates the Datalink, generates new metadata for the updated file, and transactionally updates the file with its new metadata. In this way, data files are continuously available to all users through appending and updating actions. Since updating Datalinks 25 requires the file management system to initiate backup, a "delta versioning" operation, which reduces the data needed to support backup operations, permits more efficient backup of data files and enables the continuously available files to be backed up and consistent.

30 Further, as set forth at col. 10, line 64-col. 11, line 5,

35 The DBMS interacts with the DLFM using a two-phase commit protocol to provide transaction atomicity. In this capacity, the DBMS acts as the coordinator and the DLFM acts as the subordinate. The DLFM needs some recoverable (transactional) storage for keeping persistent information about the operations (such as LinkFile, UnlinkFile, and the like) that are initiated at the DLFM by the DBMS as a result of application calls such as SQL Insert, Update, or Delete.

The reference is thus relevant to the extent that it involves the use of a two-phase commit protocol; however, as will be explained in detail in the following, it does not disclose anything

at all about optimization of two-phase commit protocols and thus cannot disclose the particular technique for optimizing two-phase commit protocol which is claimed by Applicants.

The rejection of claim 11

Claim 11 is exemplary for what Applicants are claiming:

5 11. (currently amended) A method of ensuring that a first component of a distributed system that exchanges messages belonging to a transaction with one or more other components of the distributed system is additionally aware of a state of an other component, the state being relevant to a protocol employed with the transaction and

10 the method comprising the steps practiced in the first component of:
 receiving an augmented one of the messages, the augmented message having been augmented by the other component to additionally contain information indicating the relevant state of the other component;

15 retaining the relevant state from the augmented message; and
 using the retained relevant state to optimize the protocol.

Examiner finds the “augmented message” in Burns’ link file command, which as set forth at col. 9, line 56, names a file that is to be linked to the DBMS and includes a read-only flag. The link file command is sent by the DBMS to the file manager. When the file manager receives the link file command, it sets the read only flag in the file system directory entry for the file specified in

20 the link file command to read only, as explained at col. 9, lines 63-65.

The first issue here is whether the link file command can reasonably be characterized as an “augmented message” as that term is used in Applicants’ Specification. As set forth above, augmented messages are messages that “are sent from component systems to the managing system during the transaction [and] are *augmented* with information that indicates whether the

25 component system sending the message is currently read-only with regard to the transaction.” Here, the read-only flag is not something which augments, or is added to the link file command, but is simply part of the command.

Next is the fact that the claim requires that the augmented message have been augmented by the “other component”, that is, the flow of state information is from the subordinate component to

30 the managing component. According to col. 10, line 64-col. 11, line 5 of Burns, in Burns’s system, the DBMS acts as the coordinator and the DLFM acts as the subordinate. The flow of state information in Burns is thus from the manager to the subordinate, or the reverse of what is set forth in claim 11.

a For the claim's step of "using the retained relevant state to optimize the protocol", Examiner refers Applicants to col. 9, line 65-col. 10, line 4. The cited location further explains the effect of the link file command on the file specified in the command, but discloses absolutely nothing about the use of "retained relevant state to optimize the protocol". Further, as already pointed 5 out, there is no disclosure whatever concerning optimization in Burns, as can be seen from the fact that a Lexis search on patno = (6,088,694) and optimiz! yielded no results.

Because Burns discloses neither claim 11's augmented message nor the flow of state information set forth in claim 11 nor anything whatever about optimization, Burns cannot provide a foundation for a rejection of claim 11 under 35 U.S.C. 102. As Examiner is of course 10 aware, exactly the same arguments can be made with regard to independent claims 5, 10, 22, and 26, and consequently, Burns also offers no foundation for the rejection of those claims under 35 U.S.C. 102.

The rejections of the dependent claims

The dependent claims are all patentable because they are dependent from patentable claims. In 15 addition, claims 2, 3, 6, 7, 23, 24, 27, and 28 all add details concerning the augmented messages and the use of the retained state from the augmented messages in optimizing the protocols, and are thus patentable in their own rights over reference that shows neither the augmented messages nor the use of the retained state in optimizing the protocols. Claims 4, 8, 25, and 29 set forth that the distributed system is a distributed database system. As is apparent from Burns' 20 FIG. 2 and the discussion of that figure at col. 7, line 64, col. 8, line 14, it is the *file system* that is distributed, not the database system, as required by the claim. Claims 12-21 are dependent Beauregard claims and are patentable for the reasons that the claims they are dependent from are patentable.

The rejections under 35 U.S.C. 103

25 In these rejections, Examiner combines the disclosure of Burns with that of Lampson, Lampson is being added to supply "the disclosure of an abort step in the process of two-phase commit" (Examiner's *Response to Arguments*). Since Examiner depends on Burns for all of the other limitations in the claims, the rejections can stand only if Burns in fact shows all of those limitations.

The rejections of the dependent claims

Claims 3, 7, 24, and 28 are ultimately dependent from independent claims 11, 5, 22, and 26, and consequently, the rejections depend upon those independent claims being anticipated by Burns. As set forth in detail above, the independent claims are not anticipated by Burns. Consequently, 5 Examiner has failed to establish a *prima facie* case of obviousness with regard to claims 3,7,24, and 28 and the rejection of these claims under 35 U.S.C. 103 cannot stand.

The rejections of independent claims 9, 30, and 31

These independent claims claim Applicants' optimization techniques in the particular context of the two-phase commit protocol. The limitations of all of these claims include a message that 10 has been augmented with state information indicating whether the transaction will modify the cohort and sent from the cohort to the coordinator and the coordinator's use of the state information to determine what message it sends to the cohort. As set forth above in the discussion of claim 11, Burns does not disclose these limitations; consequently, the combination of Burns and Lampson does not show all of the claimed limitations and Examiner has failed to 15 establish his *prima facie* case of obviousness with regard to claims 9, 30, and 31.

Conclusion

Applicants have amended claim 1 to correct a typographical error. The claim as amended is fully supported by the application as filed. Applicants have further traversed all of Examiner's rejections of the claims. Applicants have thereby satisfied the requirements of 37 C.F.R. 5 1.111(b) and respectfully request that Examiner continue with his examination and allow the claims as presently amended. No fees are believed to be required for the amendment. Should any be, please charge the additional fees to deposit account number 501315.

10

Respectfully submitted,



Attorney of record,
Gordon E. Nelson
57 Central St., P.O. Box 782
Rowley, MA, 01969,
Registration number 30,093
Voice: (978) 948-7632
Fax: (866) 723-0359

9/7/04

Date

15

20

Certificate of Mailing

25

I hereby certify that this correspondence is being deposited with the United States Postal Service with sufficient postage as first class mail in an envelope addressed to:

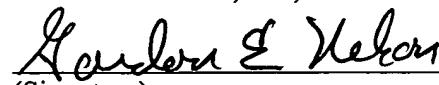
30

Assistant Commissioner for Patents
Post Office Box 1450,
Alexandria, VA 22313-1450

on 9/7/04
(Date)

35

Gordon E. Nelson, #30,093


(Signature)

40